arXiv:1403.2332v6 [stat.ME] 4 Jun 2015

A Mixture of Coalesced Generalized Hyperbolic Distributions

Cristina Tortora,* Brian C. Franczak, Ryan P. Browne
and Paul D. McNicholas

Department of Mathematics & Statistics, McMaster University.

Abstract

A mixture of coalesced generalized hyperbolic distributions (GHDs) is developed by joining a finite mixture of generalized hyperbolic distributions with a novel mixture of multiple scaled generalized hyperbolic distributions (MSGHDs). After detailing the development of the mixture of MSGHDs, which arises via implementation of a multi-dimensional weight function, the density of our coalesced distribution is developed. A parameter estimation scheme is developed using the ever-expanding class of MM algorithms and the Bayesian information criterion is used for model selection. We use our mixture of coalesced GHDs for clustering and compare them to mixtures of GHDs, mixtures of MSGHDs, and mixtures of skew-t distributions using simulated and real data.

1 Introduction

The distribution of a random variable $\mathbf{X} \in \mathbb{R}^p$ is said to be a normal variance-mean mixture if its density can be written as

$$f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \boldsymbol{\theta}) = \int_0^\infty \phi_p \left(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma}\right) f_W \left(w \mid \boldsymbol{\theta}\right) dw, \tag{1}$$

where $\phi_p \left(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma}\right)$ is the density of a p-dimensional Gaussian distribution with mean $\boldsymbol{\mu} + w\boldsymbol{\alpha}$ and covariance matrix $w\boldsymbol{\Sigma}$, and $f_W \left(w \mid \boldsymbol{\theta}\right)$, the probability density function of a univariate random variable $W > 0$, is a weight function (cf. Barndorff-Nielsen et al., 1982; Gneiting, 1997). The weight function is free to take on many forms, e.g., $f_W \left(w \mid \boldsymbol{\theta}\right)$ could

*Department of Mathematics & Statistics, McMaster University, Hamilton, Ontario, L8S 4L8, Canada. E-mail: ctortora@mcmaster.ca.

be the density of a random variable from a gamma distribution, exponential distribution, or generalized inverse Gaussian (GIG) distribution. Depending on the choice of $f_W(w \mid \boldsymbol{\theta})$, evaluating the integral in (1) can lead to a number of distinct representations for several non-Gaussian multivariate density functions (e.g., Barndorff-Nielsen, 1978; Kotz et al., 2001; Kotz and Nadarajah, 2004).

Forbes and Wraith (2014) show that a multi-dimensional weight variable $\boldsymbol{\Delta}_{\mathbf{w}} = \text{diag}(W_1, \ldots, W_p)$ can be incorporated within the density of a normal variance-mean mixture via an eigen-decomposition of the symmetric positive-definite matrix, $\boldsymbol{\Sigma}$, in (1). Formally, they set $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$, where $\boldsymbol{\Gamma}$ is a $p \times p$ matrix of eigenvectors and $\boldsymbol{\Phi}$ is a $p \times p$ diagonal matrix containing the eigenvalues of $\boldsymbol{\Sigma}$. It follows that the density of $\mathbf{X}$ becomes

$$\begin{aligned} &f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\theta}) \\ &\qquad = \int_0^\infty \cdots \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu} + \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\alpha}, \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}\boldsymbol{\Gamma}') f_{\mathbf{W}}(w_1, \ldots, w_p \mid \boldsymbol{\theta}) \, dw_1 \ldots dw_p, \end{aligned} \tag{2}$$

where $f_{\mathbf{W}}(w_1, \ldots, w_p \mid \boldsymbol{\theta}) = f_W(w_1 \mid \boldsymbol{\theta}_1) \times \cdots \times f_W(w_p \mid \boldsymbol{\theta}_p)$ is a p-variate density function such that each w_j, for $j = 1, \ldots, p$, is independent and $\phi_p(\mathbf{x} \mid \boldsymbol{\mu} + \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\alpha}, \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}\boldsymbol{\Gamma}')$ is the density of the multivariate Gaussian distribution with mean $\boldsymbol{\mu} + \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\alpha}$ and covariance matrix $\boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$. The density given in (2) adds flexibility to normal variance-mean mixtures because the parameters $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_p$ are now free to vary in each dimension, p. Using the density in (2), Forbes and Wraith (2014) derive the density of a multiple scaled multivariate-t distribution, Franczak et al. (2015) develop a multiple scaled shifted asymmetric Laplace distribution, and Wraith and Forbes (2015) derive a multiple scaled normal inverse Gaussian distribution. In each case, the respective authors use finite mixtures (see Section 2.1) of their multiple scaled distributions for cluster analysis.

In general, recognizing that a random variable $\mathbf{X}$ from some non-Gaussian distribution can be represented as a normal variance-mean mixture is advantageous when using finite mixture models for cluster analysis. Peel and McLachlan (2000), Karlis and Santourian (2009), Browne and McNicholas (2015), Franczak et al. (2014), Murray et al. (2014), and Tortora et al. (2015) all exploit the fact that $\mathbf{X}$ is a normal variance-mean mixture to derive mathematically tractable parameter estimates, via the expectation-maximization (EM) algorithm (Dempster et al., 1977), for a host of non-Gaussian distributions. These distributions would otherwise have no closed-form EM solutions.

In this paper we introduce a multiple scaled generalized hyperbolic distribution (MSGHD) and two variations of this MSGHD: the convex MSGHD (cMSGHD) and the coalesced generalized hyperbolic distributions (CGHDs). We use these distributions for model-based clustering, classification and discriminant analysis. The MSGHD parameterizes skewness, location, scale, and, unlike the generalized hyperbolic distributions (GHDs; McNeil et al., 2005), estimates the concentration and index parameters of the density in each dimension. The density of the MSGHD may not be convex, i.e., does not have convex contours; however, by adding a constraint on the index parameters, we can obtain a convex MSGHD (cMSGHD).

The CGHD merges the existing GHDs with our MSGHDs, such that its density can be the MSGHD, the GHD, or a combination of the two.

The remainder of this paper is organized as follows. In Section 2, relevant background and literature review is presented. Then, our methodology is presented (Section 3): we derive the MSGHD (Section 3.1), the MCGHDs (Section 3.2), we give explicit details concerning our parameter estimation procedure (Section 3.4), and discuss model selection (Section 3.5). In Section 3.6 we present the cMSGHD, and we then discuss our models in a classification and discriminant analysis context (Section 3.7). In Section 4, we show that the MCGHDs generally performs as well as the best performer between the mixture of GHDs (MGHDs) and the mixture of MSGHDs (MMSGHDs), and we compare our model to other approaches using simulated and real data sets. This paper concludes with a summary and discussion of future work (Section 6).

2 Background

2.1 Finite Mixture Models

Finite mixture models assume that a population is a convex combination of a finite number of probability densities. They are often utilized as either a semi-parametric alternative to nonparametric density estimation techniques or for model-based clustering and classification (cf. Titterington et al., 1985; McLachlan and Peel, 2000). Formally, a random vector $\mathbf{X}$ follows a (parametric) finite mixture distribution if, for all $\mathbf{x} \subset \mathbf{X}$, its density can be written

$$f(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g f_g(\mathbf{x} \mid \boldsymbol{\zeta}_g), \tag{3}$$

where $\pi_g > 0$, such that $\sum_{g=1}^{G} \pi_g = 1$, is the gth mixing proportion, $f_g(\mathbf{x} \mid \boldsymbol{\zeta}_g)$ is the gth component density, and $\boldsymbol{\vartheta} = (\boldsymbol{\pi}, \boldsymbol{\zeta}_1, \ldots, \boldsymbol{\zeta}_G)$ is the vector of parameters, with $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_G)$.

For model-based clustering and classification applications, the Gaussian mixture model is the most popular tool (e.g., Celeux and Govaert, 1995; Fraley and Raftery, 2002; McLachlan et al., 2003; Bouveyron et al., 2007; McNicholas and Murphy, 2008, 2010; Baek et al., 2010; Montanari and Viroli, 2011; Bouveyron and Brunet-Saumard, 2014). However, over the past few years a movement toward developing non-Gaussian finite mixtures for clustering and classification has gained momentum; resulting in an increase in available methods, including work by Karlis and Meligkotsidou (2007); Lin (2009, 2010), Browne et al. (2012), Lee and McLachlan (2013b), and Vrbik and McNicholas (2012, 2014), amongst others.

2.2 Multiple Scaled Distributions

As mentioned in Section 1, Forbes and Wraith (2014) develop a multiple scaled multivariate-t distribution. In addition, they also discuss the development of three other multiple scaled distributions: a multivariate representation of a Pearson type VII distribution (cf. Johnson et al., 1994, vol. 2, chap. 28), the so-called multivariate K model (Eltoft et al., 2006), and a multivariate normal inverse Gaussian distribution (Karlis and Santourian, 2009). They show that one representation of the multivariate-t distribution's density function arises by setting $\boldsymbol{\alpha} = \mathbf{0}$ and $f_W(w \mid \boldsymbol{\theta}) = g(w \mid \nu/2, \nu/2)$ in (1), where $g(w \mid \nu/2, \nu/2) = w^{\nu/2-1}\Gamma(\nu/2)^{-1}\exp\{-\nu w/2\}(\nu/2)^{\nu/2}$, for $w > 0$, is the density of a univariate Gamma distribution. Under this parametrization, (1) becomes

$$\begin{aligned} t_p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu) &= \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}/w)\, g(w \mid \nu/2, \nu/2)\, dw \\ &= \frac{\Gamma((\nu+p)/2)}{|\boldsymbol{\Sigma}|^{1/2}\Gamma(\nu/2)(\pi\nu)^{p/2}}\left[1 + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})/\nu\right]^{-(\nu+p)/2}, \end{aligned} \tag{4}$$

where $\phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}/w)$ is the density of a p-dimensional Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}/w$, ν is the number of degrees of freedom, and $\delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma}) = (\mathbf{x}-\boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})$ is the squared Mahalanobis distance between $\mathbf{x}$ and $\boldsymbol{\mu}$.

Multiple scaled distributions are characterized by a multivariate weight function, $g_{\mathbf{W}}(w_1, \ldots, w_p \mid \boldsymbol{\nu})$. As discussed in Section 1, by letting $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$, it follows from (2) that the density in (4) can be written

$$t_{\text{MS}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\nu}) = \int_0^\infty \cdots \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Gamma}')\, g_{\mathbf{W}}(\mathbf{w}|\boldsymbol{\nu})\, dw_1 \ldots dw_p, \tag{5}$$

where $\boldsymbol{\Delta}_{\mathbf{w}} = \text{diag}\left(w_1^{-1}, \ldots, w_p^{-1}\right)$ and the weight function

$$g(w_1, \ldots, w_p \mid \nu_j/2, \nu_j/2) = g(w_1 \mid \nu_1/2, \nu_1/2) \times \cdots \times g(w_p \mid \nu_p/2, \nu_p/2)$$

is a p-variate gamma density, where $g(w_j \mid \nu_j/2, \nu_j/2)$ is univariate gamma density function defined just above (4).

The scaled Gaussian density in (5) can be written

$$\phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Gamma}') = \prod_{j=1}^p \phi_1\left([\boldsymbol{\Gamma}'\mathbf{x}]_j \mid [\boldsymbol{\Gamma}'\boldsymbol{\mu}]_j, \Phi_j w_j^{-1}\right) = \prod_{j=1}^p \phi_1\left(\boldsymbol{\Gamma}'[\mathbf{x}-\boldsymbol{\mu}]_j \mid 0, \Phi_j w_j^{-1}\right), \tag{6}$$

where $\boldsymbol{\Gamma}'[\mathbf{x}-\boldsymbol{\mu}]_j$ is the jth element of $\boldsymbol{\Gamma}'(\mathbf{x}-\boldsymbol{\mu})$, $\phi_1(\boldsymbol{\Gamma}'[\mathbf{x}-\boldsymbol{\mu}]_j \mid 0, \Phi_j w_j^{-1})$ is the density of a univariate Gaussian distribution with mean 0 and variance $\Phi_j w_j^{-1}$, and Φ_j the jth diagonal element of the matrix $\boldsymbol{\Phi}$, i.e., the jth eigenvalue of $\boldsymbol{\Phi}$. It follows that (5) can be written

$$t_{\text{MS}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\nu}) = \prod_{j=1}^p \int_0^\infty \phi_1\left(\boldsymbol{\Gamma}'[\mathbf{x}-\boldsymbol{\mu}]_j \mid 0, \Phi_j w_j^{-1}\right) g(w_j \mid \nu_j/2, \nu_j/2)\, dw_j, \tag{7}$$

Solving the integral in (7) gives

$$t_{\text{MS}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\nu}) = \prod_{j=1}^{p} \frac{\Gamma((\nu_j+1)/2)}{\Gamma(\nu_j/2)(\Phi_j \nu_j \pi)^{1/2}} \left[1 + \frac{\left(\boldsymbol{\Gamma}' \left[\mathbf{x} - \boldsymbol{\mu}\right]_j\right)^2}{\Phi_j \nu_j}\right]^{-(\nu_j+1)/2}, \tag{8}$$

where $\Gamma(\cdot)$ is the Gamma function, Φ_j is the jth eigenvalue, $\boldsymbol{\Gamma}$ is a matrix of eigenvectors, $\boldsymbol{\mu}$ is a location parameter, and $\boldsymbol{\Gamma}' \left[\mathbf{x} - \boldsymbol{\mu}\right]_j^2 / \Phi_j$ can be regarded as the Mahalanobis distance between $\mathbf{x}$ and $\boldsymbol{\mu}$. Note that (8) is the density of a multiple scaled multivariate-t distribution.

The main difference between the traditional multivariate-t density given in (4) and the multiple scaled multivariate-t density given in (8) is that the degrees of freedom can now be parameterized separately in each dimension, p. Therefore, unlike the multivariate-t distribution, the multiple scaled representation can account for differences in tail weight in every dimension (see Forbes and Wraith, 2014, for details).

2.3 Mixtures of Generalized Hyperbolic Distributions

A generalized hyperbolic distribution arises from a normal variance-mean mixture when the weight function $f_W(w \mid \boldsymbol{\theta})$ is the density of a GIG distribution. We write $W \backsim \text{GIG}(\psi, \chi, \lambda)$ to indicate that the random variable W follows a GIG distribution with density

$$h_W(w \mid \psi, \chi, \lambda) = \frac{(\psi/\chi)^{\lambda/2} w^{\lambda-1}}{2K_\lambda(\sqrt{\psi\chi})} \exp\left\{-\frac{\psi w + \chi/w}{2}\right\}, \tag{9}$$

where $\psi, \chi \in \mathbb{R}^+$ are concentration parameters, and K_λ is the modified Bessel function of the third kind with index $\lambda \in \mathbb{R}$. McNeil et al. (2005) write the density of a random variable $\mathbf{X} \in \mathbb{R}^p$ from a generalized hyperbolic distribution as

$$\begin{aligned} f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \psi, \chi, \lambda) &= \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma})\, h_W(w \mid \psi, \chi, \lambda)\, dw \\ &= \left[\frac{\chi + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})}{\psi + \boldsymbol{\alpha}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}}\right]^{(\lambda - p/2)/2} \frac{(\psi/\chi)^{\lambda/2} K_{\lambda-p/2}\left(\sqrt{[\psi + \boldsymbol{\alpha}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}][\chi + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})]}\right)}{(2\pi)^{p/2}|\boldsymbol{\Sigma}|^{1/2} K_\lambda(\sqrt{\chi\psi}) \exp\{(\boldsymbol{\mu} - \mathbf{x})'\boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}\}}, \end{aligned} \tag{10}$$

where $\delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})$ is defined for (4), $\boldsymbol{\alpha}$ is a p-dimensional skewness parameter, $\boldsymbol{\Sigma}$ is a $p \times p$ scale matrix, and K_λ, ψ, χ, and λ are defined for (9). The random variable $\mathbf{X}$ can be generated via the relationship

$$\mathbf{X} = \boldsymbol{\mu} + W\boldsymbol{\alpha} + \sqrt{W}\mathbf{V}, \tag{11}$$

where $\mathbf{V} \backsim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and $W \backsim \text{GIG}(\psi, \chi, \lambda)$. Note that, we use $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ to represent a multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$.

Browne and McNicholas (2015) give an alternative representation of the generalized inverse Gaussian density by setting $\omega = \sqrt{\psi\chi}$ and $\eta = \sqrt{\chi/\psi}$. Formally, this gives

$$h_W\left(w \mid \omega, \eta, \lambda\right) = \frac{(w/\eta)^{\lambda-1}}{2\eta K_\lambda(\omega)} \exp\left\{-\frac{\omega}{2}\left(\frac{w}{\eta} + \frac{\eta}{w}\right)\right\}, \tag{12}$$

where $\eta > 0$ is a scale parameter, $\omega > 0$ is a concentration parameter, and K_λ and λ are as previously defined. Herein, we write $W \backsim \mathcal{I}(\omega, \eta, \lambda)$ to indicate that W follows the GIG distribution with density parameterized as in (12). The GIG distribution has a number of attractive properties. For our purposes (see Section 3.4), the most appealing of these properties are the tractability of its expected values:

$$\begin{aligned} &\mathbb{E}\left[W\right] = \eta\frac{K_{\lambda+1}\left(\omega\right)}{K_\lambda\left(\omega\right)}, \qquad \mathbb{E}\left[1/W\right] = \frac{1}{\eta}\frac{K_{\lambda-1}\left(\omega\right)}{K_\lambda\left(\omega\right)} = \frac{1}{\eta}\frac{K_{\lambda+1}\left(\omega\right)}{K_\lambda\left(\omega\right)} - \frac{2\lambda}{\omega\eta}, \\ &\mathbb{E}\left[\log W\right] = \log\eta + \frac{\partial}{\partial\lambda}\log K_\lambda\left(\omega\right) \end{aligned} \tag{13}$$

where K_λ is as previously defined (cf. Jørgensen, 1982).

To ensure identifiability, the density given in (10) requires the constraint $|\boldsymbol{\Sigma}| = 1$; however, because this would be detrimental to model-based clustering and classification applications, Browne and McNicholas (2015) modify the stochastic relationship in (11) such that

$$\mathbf{X} = \boldsymbol{\mu} + W\eta\boldsymbol{\alpha} + \sqrt{W\eta}\mathbf{V} = \boldsymbol{\mu} + W\boldsymbol{\beta} + \sqrt{W}\mathbf{V}, \tag{14}$$

where $\boldsymbol{\beta} = \eta\boldsymbol{\alpha}$, $\mathbf{V}$ is as defined for (11), and set the scale parameter $\eta = 1$ such that $W \backsim \mathcal{I}\left(\omega, 1, \lambda\right)$. The effect of the constraint $\eta = 1$ is to set $\psi = \chi$ so that there is only one concentration parameter, i.e., ω.

Under this parametrization, the density of the generalized hyperbolic distribution is

$$\begin{aligned} f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\beta}, \omega, \lambda) = &\left[\frac{\omega + \delta\left(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma}\right)}{\omega + \boldsymbol{\beta}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\beta}}\right]^{(\lambda-p/2)/2} \\ &\times \frac{K_{\lambda-p/2}\left(\sqrt{\left[\omega + \boldsymbol{\beta}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\beta}\right]\left[\omega + \delta\left(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma}\right)\right]}\right)}{(2\pi)^{p/2}\left|\boldsymbol{\Sigma}\right|^{1/2} K_\lambda\left(\omega\right)\exp\left\{-\left(\mathbf{x} - \boldsymbol{\mu}\right)'\boldsymbol{\Sigma}^{-1}\boldsymbol{\beta}\right\}}, \end{aligned} \tag{15}$$

where $\boldsymbol{\beta}$, ω, and λ are defined as before, and $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are location and scale parameters, respectively. It follows that the density of a mixture of GHDs is given by

$$f_{\text{MGH}}(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, \boldsymbol{\beta}_g, \omega_g, \lambda_g), \tag{16}$$

where $f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, \boldsymbol{\beta}_g, \omega_g, \lambda_g)$ is the density of the generalized hyperbolic distribution given in (15) and π_g are the mixing proportions as defined for (3).

3 Methodology

3.1 A multiple scaled generalized hyperbolic distribution

To extend the generalized hyperbolic distribution to its multiple scaled version, we first note that the relationship given in (14) can be transformed via an eigen-decomposition of the matrix $\boldsymbol{\Sigma}$ and the introduction of a multi-dimensional weight variable. Specifically, we can write that a random variable $\mathbf{X}$ from a MSGHD can be generated via the relationship

$$\mathbf{X} = \boldsymbol{\Gamma}\boldsymbol{\mu} + \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\beta} + \boldsymbol{\Gamma}\mathbf{V}, \tag{17}$$

where $\mathbf{V} \backsim \mathcal{N}(\mathbf{0}, \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi})$ and $\boldsymbol{\Delta}_{\mathbf{w}} = \text{diag}(w_1, \ldots, w_p)$. Therefore, it follows that $\mathbf{X} \mid \mathbf{w} \backsim \mathcal{N}(\boldsymbol{\Gamma}\boldsymbol{\mu} + \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\beta}, \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}\boldsymbol{\Gamma}')$ and that the density of $\mathbf{X}$ can be written

$$f_{\text{MSGH}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}) = \int_0^\infty \cdots \int_0^\infty \phi_p(\boldsymbol{\Gamma}'\mathbf{x} - \boldsymbol{\mu} - \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\beta} \mid \mathbf{0}, \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi})\, h_{\mathbf{w}}(w_1, \ldots, w_p \mid \boldsymbol{\omega}, \mathbf{1}, \boldsymbol{\lambda}) d\mathbf{w}, \tag{18}$$

where $\phi_p(\boldsymbol{\Gamma}'\mathbf{x} - \boldsymbol{\mu} - \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\beta} \mid \mathbf{0}, \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi})$ is the density of a multivariate Gaussian distribution with mean $\mathbf{0}$ and covariance $\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}$, and $h_{\mathbf{w}}(w_1, \ldots, w_p \mid \boldsymbol{\omega}, \mathbf{1}, \boldsymbol{\lambda})$ is the density of a p unidimensional generalized inverse Gaussian distributions given in equation (12) with $\eta = 1$. Therefore, following the derivations that lead to (8), we find that the density of a random variable $\mathbf{X}$ from a MSGHD is given by

$$\begin{aligned} f_{\text{MSGH}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}) &= \prod_{j=1}^{P} \int_0^\infty \phi_1\left([\boldsymbol{\Gamma}'\mathbf{x} - \boldsymbol{\mu} - \boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\beta}]_j \mid 0, \Phi_j w_j\right) h_W(w_j \mid \omega_j, 1, \lambda_j)\, dw_j \\ &= \prod_{j=1}^{P} \left[\frac{\omega_j + \Phi_j^{-1}\left([\boldsymbol{\Gamma}'\mathbf{x}]_j - \boldsymbol{\mu}_j\right)^2}{\omega_j + \boldsymbol{\beta}_j^2 {\Phi_j}^{-1}}\right]^{\frac{\lambda_j - \frac{1}{2}}{2}} \frac{K_{\lambda_j - \frac{1}{2}}\left(\sqrt{[\omega_j + \boldsymbol{\beta}_j^2 {\Phi_j}^{-1}]\left[\omega_j + \Phi_j^{-1}\left([\boldsymbol{\Gamma}'\mathbf{x}]_j - \boldsymbol{\mu}_j\right)^2\right]}\right)}{(2\pi)^{\frac{1}{2}} \Phi_j^{\frac{1}{2}} K_{\lambda_j}(\omega_j) \exp\left\{\left([\boldsymbol{\Gamma}'\mathbf{x}]_j - \boldsymbol{\mu}_j\right)\boldsymbol{\beta}_j\right\}}, \end{aligned} \tag{19}$$

where $[\boldsymbol{\Gamma}'\mathbf{x}]_j$ is the j^{th} element of the vector $\boldsymbol{\Gamma}'\mathbf{x}$, $\boldsymbol{\mu}_j$ is the j^{th} element of the location parameter $\boldsymbol{\mu}$, $\boldsymbol{\beta}_j$ is the j^{th} element of the skewness parameter $\boldsymbol{\beta}$, $\boldsymbol{\Gamma}$ is a $p \times p$ matrix of eigenvectors, Φ_j is the jth eigenvalue of the diagonal matrix $\boldsymbol{\Phi}$, $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_p)'$ controls the concentration in each dimension p, and $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_p)'$ is a p-dimensional index parameter.

3.2 A coalesced generalized hyperbolic distribution

The GHD is not a special case of the MSGHD under any parameterization when $p > 1$. Accordingly, we propose a novel coalesced distribution that contains both the GHD and

MSGHD as special cases. Our coalesced distribution arises through the introduction of a random vector

$$\mathbf{R} = U\mathbf{X} + (1-U)\mathbf{S}, \tag{20}$$

where $\mathbf{S} \backsim \mathcal{MSGHD}\left(\boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}\right)$, $\mathbf{X} = \boldsymbol{\Gamma}\mathbf{Y}$, where $\mathbf{Y} \backsim \mathcal{GHD}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\beta}, \omega_0, \lambda_0\right)$, with $\boldsymbol{\Sigma}$ decomposed into the constituent elements of an eigen-decomposition, i.e., $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$, and U is an indicator variable such that when $U = 1$ R follows a GHD and when $U = 0$ R follows a MSGHD. It follows that $\mathbf{X} = \boldsymbol{\Gamma}\boldsymbol{\mu} + W\boldsymbol{\Gamma}\boldsymbol{\beta} + \sqrt{W}\boldsymbol{\Gamma}\mathbf{V}$, where $\boldsymbol{\Gamma}\mathbf{V} \backsim \mathcal{N}_p\left(\mathbf{0}, \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\right)$, cf. (14), $\mathbf{S} = \boldsymbol{\Gamma}\boldsymbol{\mu} + \boldsymbol{\Gamma}\boldsymbol{\beta}\boldsymbol{\Delta}_{\mathbf{w}} + \boldsymbol{\Gamma}\mathbf{A}$, where $\boldsymbol{\Gamma}\mathbf{A} \backsim \mathcal{N}_p\left(\mathbf{0}, \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\right)$, and the density of $\mathbf{R}$ can be written

$$\begin{aligned} f_{\text{CGH}}\left(\mathbf{r} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi\right) = \\ \varpi f_{\text{GH}}\left(\mathbf{r} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}', \boldsymbol{\beta}, \omega_0, \lambda_0\right) + (1-\varpi) f_{\text{MSGH}}\left(\mathbf{r} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}\right), \end{aligned} \tag{21}$$

where f_{GH} is the density of a generalized hyperbolic random variable, given in (15), f_{MSGH} is the density of the MSGHD random variable given in (19), and $\varpi \in [0, 1]$ is the inner mixing proportion, defined for (20). It follows that the random vector $\mathbf{R}$ will be distributed generalized hyperbolic if $\varpi = 1$ and will be distributed multiple scaled generalized hyperbolic if $\varpi = 0$. The parameters $\boldsymbol{\mu}$, $\boldsymbol{\beta}$, $\boldsymbol{\Gamma}$, and $\boldsymbol{\Phi}$ are the same for both densities, the parameters ω_0 and λ_0 are univariate values unique to the GHD, and the p-dimensional parameters $\boldsymbol{\omega}$ and $\boldsymbol{\lambda}$ are unique to the MSGHD. Herein, we write $\mathbf{R} \backsim \text{CGHD}(\boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi)$ to indicate that the random vector $\mathbf{R}$ follows a CGHD and we indicate with $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$.

3.3 Mixtures of CGHDs

Herein, we use mixtures of CGHDs (MCGHDs) for model-based clustering and classification. The mixtures of CGHDs has density

$$f_{\text{MCGH}}(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g f_{\text{CGH}}\left(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\beta}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g, \omega_{0g}, \lambda_{0g}, \varpi_g\right), \tag{22}$$

where π_g are the mixing proportions and f_{CGH} is the density given in (21). It follows from Section 3.2 that the MCGHDs contains both the MGHDs, whose density is given in (16), and mixtures of MSGHDs (MMSGHDs) whose density is given by

$$f_{\text{MMSGH}}(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g f_{\text{MSGH}}\left(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\beta}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g\right), \tag{23}$$

where π_g are the mixing proportions and f_{MSGH} is the density given in (19).

3.4 Parameter estimation

To estimate the parameters of the MCGHD, we use an EM algorithm. The EM algorithm iterates between two steps, an E-step and a M-step, until convergence. On each E-step,

the expected value of the complete-data log-likelihood, $\mathcal{Q}$, is calculated and, on each M-step, this expected value is maximized with respect to the model parameters. Within our EM algorithm, we draw on a class of algorithms known as MM algorithms (Ortega and Rheinboldt, 1970; Hunter and Lange, 2000); MM stands for 'minorize-maximize' or 'majorize-minimize,' depending on the purpose of the algorithm. In our M-step, we increase $\mathcal{Q}$ rather than maximizing it; accordingly, ours is formally a generalized EM (GEM) algorithm.

For our MCGHDs, there are four sources of missing data: the latent variable W_{0ig}, the multi-dimensional weight variable $\boldsymbol{\Delta}_{\mathbf{w}ig} = \text{diag}(W_{i1g}, \ldots, W_{ipg})$, the group component indicator labels Z_{ig}, and inner component labels U_{ig}, for $i = 1, \ldots, n$ and $g = 1, \ldots, G$. For each observation i, $Z_{ig} = 1$ if observation i is in component g and $Z_{ig} = 0$ otherwise. Similarly, for each observation i, $U_{ig} = 1$ if observation i, in component g, is distributed generalized hyperbolic and $U_{ig} = 0$ if observation i, in component g, is distributed multiple scaled generalized hyperbolic. It follows that the complete-data log-likelihood for the MCGHDs is given by

$$\begin{aligned} l_c\left(\boldsymbol{\vartheta}_g\right) = \sum_{i=1}^{n}\sum_{g=1}^{G} \Big\{ & z_{ig}\log\pi_g + z_{ig}u_{ig}\log\varpi_g + z_{ig}(1-u_{ig})\log(1-\varpi_g) \\ & + z_{ig}u_{ig}\log h_W\left(w_{0ig}|\omega_{0g},1,\lambda_{0g}\right) + z_{ig}(1-u_{ig})\sum_{j=1}^{p} h_W\left(w_{ijg}|\omega_{jg},1,\lambda_{jg}\right) \\ & + z_{ig}u_{ig}\log\phi_p\left(\boldsymbol{\Gamma}_g'\mathbf{x}_i|\boldsymbol{\mu}_g + w_{0ig}\boldsymbol{\beta}_g, w_{0ig}\boldsymbol{\Phi}\right) \\ & + z_{ig}(1-u_{ig})\sum_{j=1}^{p}\log\phi_1\left([\boldsymbol{\Gamma}_g'\mathbf{x}_i]_j|\mu_{jg} + w_{ijg}\beta_{jg}, \omega_{jg}\phi_{jg}\right)\Big\} \end{aligned} \tag{24}$$

where $[\boldsymbol{\Gamma}_g'\mathbf{x}_i]_j$ is the jth element of the matrix $\boldsymbol{\Gamma}_g'\mathbf{x}_i$, $\phi_p(\cdot)$ is a p-dimensional Gaussian density, and $\phi_1(\cdot)$ is a univariate Gaussian density function. On the E-step, the expected value of the complete-data log-likelihood, $\mathcal{Q}$, is computed by replacing the sufficient statistics of the missing data by their expected values. For each component indicator label, z_{ig}, and inner component label, u_{ig}, for $i = 1, \ldots, n$ and $g = 1, \ldots, G$, we require the following expectations:

$$\mathbb{E}\left[Z_{ig} \mid \mathbf{x}_i\right] = \frac{\pi_g f_{\text{CGH}}\left(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\beta}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g, \omega_{0g}, \lambda_{0g}, \varpi_g\right)}{\sum_{h=1}^{G}\pi_h f_{\text{CGH}}\left(\mathbf{x} \mid \boldsymbol{\mu}_h, \boldsymbol{\Gamma}_h, \boldsymbol{\Phi}_h, \boldsymbol{\beta}_h, \boldsymbol{\omega}_h, \boldsymbol{\lambda}_h, \omega_{0h}, \lambda_{0h}, \varpi_h\right)} =: \hat{z}_{ig} \tag{25}$$

and

$$\begin{aligned} &\mathbb{E}\left[U_{ig} \mid \mathbf{x}_i, Z_{ig} = 1\right] = \\ &\frac{\varpi_g f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g', \boldsymbol{\beta}_g, \omega_{0g}, \lambda_{0g})}{\varpi_g f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g', \boldsymbol{\beta}_g, \omega_{0g}, \lambda_{0g}) + (1-\varpi_g) f_{\text{MSGH}}\left(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\beta}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g\right)} =: \hat{u}_{ig}, \end{aligned} \tag{26}$$

where f_{CGH} is given in (21), f_{GH} is given in (15) and f_{MSGH} is given in (19).

For the latent variable W_{0ig}, we use the expected value given in Browne and McNicholas (2015). The authors show that, given the density in (15), the random variable $W_{0ig} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 1 \backsim \mathcal{GIG}\left(\omega_{0g} + \boldsymbol{\beta}_g'(\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g')^{-1}\boldsymbol{\beta}_g, \omega_{0g} + \delta(\mathbf{x}_i, \boldsymbol{\mu}_g \mid \boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g'), \lambda_{0g} - p/2\right)$. For the MCGHDs, the maximization of $\mathcal{Q}$ requires the expected values of W_{0ig}, W_{0ig}^{-1} and $\log W_{0ig}$, i.e.,

$$\begin{aligned}
\mathbb{E}[W_{0ig} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 1] &= \sqrt{\frac{e_{ig}}{d_g}}\frac{K_{\lambda_{0g}-p/2+1}\left(\sqrt{d_g e_{ig}}\right)}{K_{\lambda_{0g}-p/2}\left(\sqrt{d_g e_{ig}}\right)} =: a_{ig},\\
\mathbb{E}[W_{0ig}^{-1} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 1] &= \sqrt{\frac{d_g}{e_{ig}}}\frac{K_{\lambda_{0g}-p/2+1}\left(\sqrt{d_g e_{ig}}\right)}{K_{\lambda_{0g}-p/2}\left(\sqrt{d_g e_{ig}}\right)} - \frac{2\lambda_{0g}-p}{e_{ig}} =: b_{ig},\\
\mathbb{E}[\log W_{0ig} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 1] &= \log\sqrt{\frac{e_{ig}}{d_g}} + \frac{\partial}{\partial v}\log\left\{K_v\left(\sqrt{d_g e_{ig}}\right)\right\}\Big|_{v=\lambda_{0g}-p/2} =: c_{ig},
\end{aligned}$$

where $d_g = \omega_{0g} + \boldsymbol{\beta}_g'(\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g')^{-1}\boldsymbol{\beta}_g$ and $e_{ig} = \omega_{0g} + \delta(\mathbf{x}_i, \boldsymbol{\mu}_g \mid \boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}_g')$.

The maximization of $\mathcal{Q}$ also requires the expected values of the multidimensional weight variables $\boldsymbol{\Delta}_{\mathbf{w}ig}$, $\boldsymbol{\Delta}_{\mathbf{w}ig}^{-1}$, and $\log\boldsymbol{\Delta}_{\mathbf{w}ig}$. Given the density in (19), it follows that $W_{ijg}|\mathbf{x}_i, Z_{ig} = 1, U_{ig} = 0 \backsim \mathcal{GIG}\left(\omega_{jg} + \beta_{jg}^2\Phi_{jg}^{-1}, \omega_{jg} + ([\mathbf{x}_i - \boldsymbol{\mu}_g]_j)^2/\phi_{jg}, \lambda_{jg} - 1/2\right)$. As such, we replace each multidimensional weight variable with its expected value, i.e., one of $\boldsymbol{E}_{1ig} = \text{diag}\{E_{1i1g}, \ldots, E_{1ipg}\}$, $\boldsymbol{E}_{2ig} = \text{diag}\{E_{2i1g}, \ldots, E_{2ipg}\}$, and $\boldsymbol{E}_{3ig} = \text{diag}\{E_{3i1g}, \ldots, E_{3ipg}\}$, where

$$\begin{aligned}
\mathbb{E}[W_{ijg} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 0] =& \sqrt{\frac{\bar{e}_{ijg}}{\bar{d}_{jg}}}\frac{K_{\lambda_{jg}+1/2}\left(\sqrt{\bar{d}_{jg}\bar{e}_{ijg}}\right)}{K_{\lambda_{jg}-1/2}\left(\sqrt{\bar{d}_{jg}\bar{e}_{ijg}}\right)} =: E_{1ijg},\\
\mathbb{E}[W_{ijg}^{-1} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 0] =& \sqrt{\frac{\bar{d}_{jg}}{\bar{e}_{ijg}}}\frac{K_{\lambda_{jg}+1/2}\left(\sqrt{\bar{d}_{jg}\bar{e}_{ijg}}\right)}{K_{\lambda_{jg}-1/2}\left(\sqrt{\bar{d}_{jg}\bar{e}_{ijg}}\right)} - \frac{2\lambda_{jg}-1}{\bar{e}_{ijg}} =: E_{2ijg},\\
\mathbb{E}[\log W_{ijg} \mid \mathbf{x}_i, Z_{ig} = 1, U_{ig} = 0] =& \log\sqrt{\frac{\bar{e}_{ijg}}{\bar{d}_{jg}}} + \frac{\partial}{\partial v}\log\left\{K_v\left(\sqrt{\bar{d}_{jg}\bar{e}_{ijg}}\right)\right\}\Big|_{v=\lambda_{jg}-1/2}\\
&=: E_{3ijg},
\end{aligned} \tag{27}$$

$\bar{d}_{jg} = \omega_{jg} + \beta_{jg}^2\Phi_{jg}^{-1}$ and $\bar{e}_{ijg} = \omega_{jg} + ([x_i - \mu_g]_j)^2/\phi_{jg}$. Furthermore, herein we let $n_g = \sum_{i=1}^n \hat{z}_{ig}$, $A_g = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}a_{ig}$, $B_g = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}b_{ig}$, $C_g = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}c_{ig}$, $\bar{E}_{1jg} = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}E_{1ijg}$, $\bar{E}_{2jg} = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}E_{2ijg}$, and $\bar{E}_{3jg} = (1/n_g)\sum_{i=1}^n \hat{z}_{ig}E_{3ijg}$.

On the M-step, we maximize the expected value of the complete-data log-likelihood with respect to the model parameters. The mixing proportions and inner mixing proportions are updated via $\hat{\pi}_g = n_g/n$ and $\hat{\varpi}_g = \sum_{i=1}^n \hat{u}_{ig}\hat{z}_{ig}/n_g$, respectively. The elements of the location parameter $\boldsymbol{\mu}_g$ and skewness parameter $\boldsymbol{\beta}_g$ are replaced with

$$\hat{\mu}_{jg} = \frac{\sum_{i=1}^n \hat{z}_{ig}[\boldsymbol{\Gamma}_g'\mathbf{x}_i]_j(\bar{s}_{1jg}s_{2ijg} - 1)}{\sum_{i=1}^n \hat{z}_{ig}(\bar{s}_{1jg}s_{2ijg} - 1)} \quad \text{and} \quad \hat{\beta}_{jg} = \frac{\sum_{i=1}^n \hat{z}_{ig}[\boldsymbol{\Gamma}_g'\mathbf{x}_i]_j(\bar{s}_{2jg} - s_{2ijg})}{\sum_{i=1}^n \hat{z}_{ig}(\bar{s}_{1jg}s_{2ijg} - 1)},$$

respectively, where $[\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j$ is the j^{th} element of the matrix $\boldsymbol{\Gamma}'_g \mathbf{x}_i$, $s_{1ijg} = \hat{u}_{ig} a_{ig} + (1 - \hat{u}_{ig}) E_{1ijg}$, $s_{2ijg} = \hat{u}_{ig} b_{ig} + (1 - \hat{u}_{ig}) E_{2ijg}$, $\bar{s}_{1jg} = 1/n_g \sum_{i=1}^{n} \hat{z}_{ig} s_{1ijg}$,$\bar{s}_{2jg} = 1/n_g \sum_{i=1}^{n} \hat{z}_{ig} s_{2ijg}$. The diagonal elements of the matrix $\boldsymbol{\Phi}_g$ are updated using

$$\begin{aligned}\hat{\phi}_{jg} &= \frac{1}{n_g} \sum_{i=1}^{n} \Big\{ \hat{z}_{ig} \hat{u}_{ig} \left[b_{ig} \left([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg} \right)^2 - 2 \left([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg} \right) \hat{\beta}_{jg} + a_{ig} \hat{\beta}_{jg}^2 \right] \\ &\qquad + \hat{z}_{ig} (1 - \hat{u}_{ig}) \left[E_{2ijg} \left([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg} \right)^2 - 2 \left([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg} \right) \hat{\beta}_{jg} + E_{1ijg} \hat{\beta}_{jg}^2 \right] \Big\}.\end{aligned}$$

To update the component eigenvector matrices $\boldsymbol{\Gamma}_g$, we wish to minimize the objective function

$$f(\boldsymbol{\Gamma}_g) = -\frac{1}{2} \operatorname{tr} \left\{ \hat{z}_{ig} \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{V}_{ig} \boldsymbol{\Gamma}_g \mathbf{x}_i \mathbf{x}_i \boldsymbol{\Gamma}'_g \right\} + \operatorname{tr} \left\{ \hat{z}_{ig} \mathbf{x}_i \left(\boldsymbol{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{\boldsymbol{\beta}}_g \right)' \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{\Gamma}_g \right\} + C \tag{28}$$

with respect to $\boldsymbol{\Gamma}_g$, where $\boldsymbol{V}_{ig} = \hat{u}_{ig} b_{ig} \mathbf{I}_p + (1 - \hat{u}_{ig}) \mathbf{E}_{2ig}$. We employ an optimization routine that uses two simpler majorization-minimization algorithms. Our optimization routine exploits the convexity of the objective function in (28), providing a computationally stable algorithm for estimating $\boldsymbol{\Gamma}_g$. Specifically, we follow Kiers (2002) and Browne and McNicholas (2014) and use the surrogate function

$$f(\boldsymbol{\Gamma}_g) \leq C + \sum_{i=1}^{n} \operatorname{tr}\{\mathbf{F}_r \boldsymbol{\Gamma}_g\}, \tag{29}$$

where C is a constant that does not depend on $\boldsymbol{\Gamma}_g$, $r \in \{1, 2\}$ is an index, and the matrices $\mathbf{F}_r$ are defined in (30) and (31).

Therefore, on each M-step, we calculate either

$$\mathbf{F}_{1g} = \sum_{i=1}^{n} \left[-\hat{z}_{ig} \mathbf{x}_i \left(\boldsymbol{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{\boldsymbol{\beta}}_g \right)' \hat{\boldsymbol{\Phi}}_g^{-1} + \hat{z}_{ig} \mathbf{x}_i \mathbf{x}'_i \boldsymbol{\Gamma}'_g \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{V}_{ig} - \hat{z}_{ig} \alpha_{1ig} \mathbf{x}_i \mathbf{x}'_i \boldsymbol{\Gamma}'_g \right] \tag{30}$$

or

$$\mathbf{F}_{2g} = \sum_{i=1}^{n} \left[-\hat{z}_{ig} \mathbf{x}_i \left(\boldsymbol{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{z}_{ig} \hat{\boldsymbol{\beta}}_g \right)' \hat{\boldsymbol{\Phi}}_g^{-1} + \hat{z}_{ig} \mathbf{x}_i \mathbf{x}'_i \boldsymbol{\Gamma}'_g \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{V}_{ig} - \hat{z}_{ig} \alpha_{2ig} \boldsymbol{V}_{ig} \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{\Gamma}'_g \right], \tag{31}$$

where α_{1ig} is the largest eigenvalue of the diagonal matrix $\boldsymbol{\Phi}_g^{-1} \boldsymbol{V}_{ig}$, and α_{2ig} is equal to $\hat{z}_{ig} \mathbf{x}'_i \mathbf{x}_i$, which is the largest eigenvalue of the rank-1 matrix $\hat{z}_{ig} \mathbf{x}_i \mathbf{x}'_i$. Following this, we compute the singular value decomposition of $\mathbf{F}_r$ given by

$$\mathbf{F}_r = \mathbf{PBR}'.$$

It follows that our update for $\boldsymbol{\Gamma}_g$ is given by

$$\hat{\boldsymbol{\Gamma}}_g = \mathbf{RP}'.$$

The p-dimensional concentration and index parameters, i.e., $\boldsymbol{\omega}_g$ and $\boldsymbol{\lambda}_g$, are estimated by maximizing the function

$$q_{jg}(\omega_{jg}, \lambda_{jg}) = -\log K_{\lambda_{jg}}(\omega_{jg}) + (\lambda_{jg} - 1)\bar{E}_{3jg} - \frac{\omega_{jg}}{2}(\bar{E}_{1jg} + \bar{E}_{2jg}). \tag{32}$$

This leads to

$$\hat{\lambda}_{jg} = \bar{E}_{3jg}\lambda_{jg}\left[\frac{\partial}{\partial v}\log K_v(\omega_{jg})\bigg|_{v=\lambda_{jg}}\right]^{-1}$$

and

$$\hat{\omega}_{jg} = \omega_{jg} - \left[\frac{\partial}{\partial v}q_{jg}(v, \lambda_{jg})\bigg|_{v=\omega_{jg}}\right]\left[\frac{\partial^2}{\partial v^2}q_{jg}(v, \lambda_{jg})\bigg|_{v=\omega_{jg}}\right]^{-1}.$$

The univariate parameters ω_{0g} and λ_{0g} are estimated by maximizing the function

$$q_{0g}(\omega_{0g}, \lambda_{0g}) = -\log(K_{\lambda_{0g}}(\omega_{0g})) + (\lambda_{0g} - 1)C_g - \frac{\omega_{0g}}{2}(A_g + B_g), \tag{33}$$

giving

$$\hat{\lambda}_{0g} = C_g\lambda_{0g}\left[\frac{\partial}{\partial v}\log K_v(\omega_{0g})\bigg|_{v=\lambda_{0g}}\right]^{-1}$$

and

$$\hat{\omega}_{0g} = \omega_{0g} - \left[\frac{\partial}{\partial v}q_{0g}(v, \lambda_{0g})\bigg|_{v=\omega_{0g}}\right]\left[\frac{\partial^2}{\partial v^2}q_{0g}(v, \lambda_{0g})\bigg|_{v=\omega_{0g}}\right]^{-1}.$$

Our GEM algorithm is iterated until convergence, which is determined using Aitken acceleration (Aitken, 1926). Formally, Aitken acceleration is given by

$$a^{(k)} = \frac{l^{(k+1)} - l^{(k)}}{l^{(k)} - l^{(k-1)}},$$

where $l^{(k)}$ is the value of the log-likelihood at the iteration k and

$$l_\infty^{(k+1)} = l^{(k)} + \frac{1}{1 - a^{(k)}}\left(l^{(k+1)} - l^{(k)}\right),$$

is an asymptotic estimate of the log-likelihood on iteration $k+1$. The algorithm can be considered to have converged when $l_\infty^{(k)} - l^{(k)} < \epsilon$, provided this difference is positive (Böhning et al., 1994; Lindsay, 1995). Herein, we let $\epsilon = 0.01$. When the algorithm converges we compute the maximum *a posteriori* (MAP) classification values using the posterior $\hat{z}_{ig}$, where $\text{MAP}\{\hat{z}_{ig}\} = 1$ if $\max_h\{\hat{z}_{ih}\}$ occurs in component $h = g$, and $\text{MAP}\{\hat{z}_{ig}\} = 0$ otherwise.

3.5 Model selection

In many clustering applications, the number of components, G, is unknown. Herein, we choose the number of components via the Bayesian information criterion (BIC; Schwarz, 1978). The BIC is defined as

$$\text{BIC} = 2l(\mathbf{x}_1, \ldots, \mathbf{x}_n \mid \hat{\boldsymbol{\vartheta}}) - \rho \log n,$$

where $l(\mathbf{x}_1, \ldots, \mathbf{x}_n \mid \hat{\boldsymbol{\vartheta}})$ is the maximized log-likelihood, $\hat{\boldsymbol{\vartheta}}$ is the vector of parameters that maximize the log-likelihood, ρ is the number of free parameters, and n is the number of observations. Arguments supporting the use of the BIC for model selection in this context are given by Campbell et al. (1997) and Dasgupta and Raftery (1998).

3.6 Convex mixture of multiple scaled generalized hyperbolic distributions

The MSGHD is more flexible than the GHD. However, like the multiple scaled multivariate-t distribution of Forbes and Wraith (2014), the MSGHD can have components that are not convex. In some situations, a mixture of MSGHDs that ensures convex components can be more suitable for clustering. Consider the data in Figure 1.

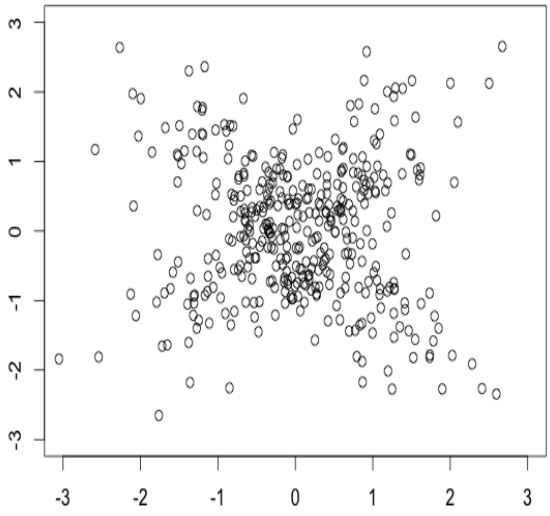

Figure 1: Data simulated from a two-component Gaussian mixture model.

How many clusters are in Figure 1? The most plausible answer is two overlapping clusters: one with positive correlation between the variables and another with negative correlation

between the variables. There may also be an argument for four or five. However, when the MMSGHD is fitted to these data for $G = 1, \ldots, 5$, the BIC selects a $G = 1$ component model (Figure 2). On the other hand, if we force the MSGHD to be convex, fitting the corresponding mixture of convex multiple scaled generalized hyperbolic distributions (McMSGHD) results in a $G = 2$ component model for these data (Figure 2). Specifically, the constraint $\lambda_j > 1$ insures convex components, i.e., if each component is associated with a cluster, then convex clusters are ensured. Formally, this amounts to insuring that the MSGHD is quasi-convex.

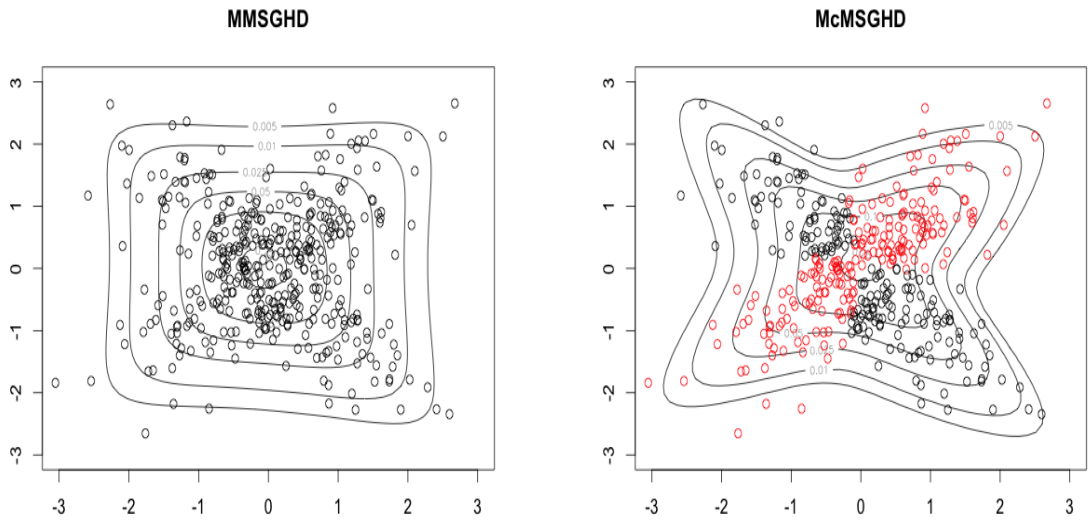

Figure 2: Contour plots for model-based clustering results on data simulated from a Gaussian mixture, using the MMSGHDs and McMSGHDs models, respectively, where colour denotes predicted classifications.

The general point here is that if convexity is not enforced, then the MSGHD can give components that contain multiple clusters. While it is easy to spot this in two dimensions, e.g., Figure 3, this phenomenon may go unrecognized in higher dimensions, possibly resulting in greatly misleading results. Of course, the issue of non-convex clusters does not arise with most model-based approaches; however, when multiple scaled mixtures are considered, the issue can crop up. Figure 3 shows some more examples to illustrate situations where the McMSGHD gives sensible clustering results and the MMSGHD gives poor results because one of the components is concave, i.e., quasi-convex. Like Figure 2, the data in Figure 3 are generated from Gaussian mixtures. The contours in Figure 3 represent the respective mixture densities; however, looking at the colours, which indicate MAP classifications, the problem becomes apparent. Specifically, in the first row of Figure 3, the MMSGHD has one convex component (associated with the red points) that does not have a sensible interpretation as a cluster. In the second row, the v-shaped MSGHD component (associated with the blue points) is certainly not a sensible cluster, while the green and red MSGHD components are also difficult to interpret as clusters. On the other hand, the McMSGHD gives sensible

clusters in both cases (Figure 3).

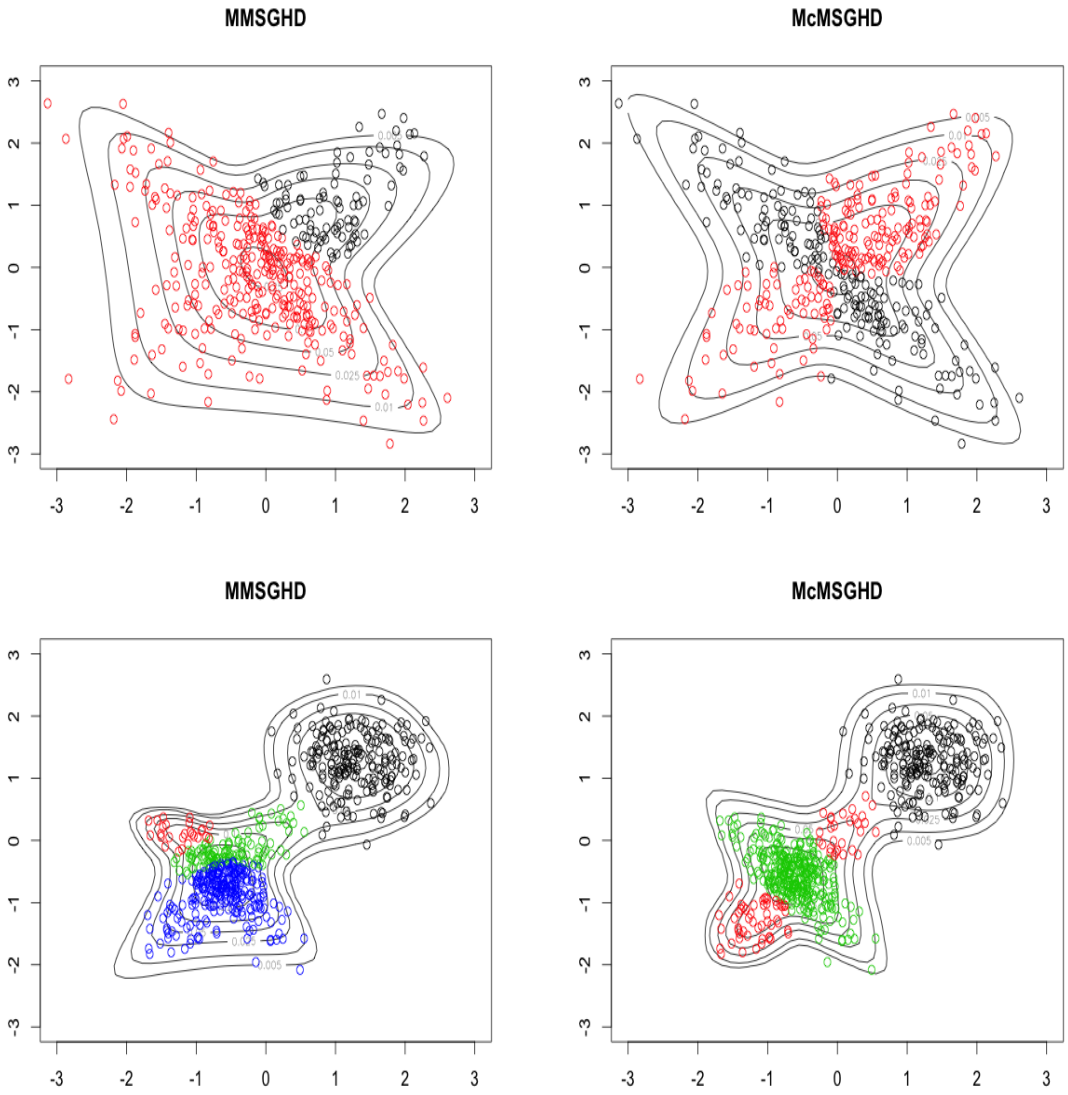

Figure 3: Contour plots for model-based clustering results on data simulated from a Gaussian mixture, using the MMSGHDs and McMSGHDs models, respectively, where colour denotes predicted classifications.

The McMSGHD is not developed with the intention of replacing the MSGHD, but rather to complement it. In a higher dimensional application, where visualization is difficult or impossible, situations where the selected MSGHD has fewer components than the selected McMSGHD will deserve special attention. Of course, this is not to say that the selected McMSGHD will always have more components in situations where the MSGHD has too few, but rather that it will help to avoid the sort of situations depicted in Figures 2 and 3.

3.7 Classification and Discriminant analysis

Model-based clustering can be seen as a special case of model-based classification. Suppose that, of the n units, k are labelled and $n-k$ are unlabelled. Model-based classification classifies the $n-k$ samples within a joint likelihood framework. Ordering the n samples so that the first k are labelled, the model-based classification likelihood for the MCGHD is

$$L(\mathbf{x} \mid \boldsymbol{\vartheta}) = \prod_{i=1}^{k} \prod_{g=1}^{G} \pi_g f_{\text{CGH}}(\mathbf{x}_i \mid \boldsymbol{\zeta}_g) \prod_{j=k+1}^{n} \sum_{h=1}^{H} \pi_h f_{\text{CGH}}(\mathbf{x}_j \mid \boldsymbol{\zeta}_h), \tag{34}$$

where $H \geq G$. Note that it is often assumed that $H = G$. Model-based classification likelihoods for the MMSGHD and McMSGHD are analogous.

In model-based discriminant analysis, the data set is divided into a training set $(\mathbf{x}_1, \ldots, \mathbf{x}_k)$ and a test set $(\mathbf{x}_{k+1}, \ldots, \mathbf{x}_n)$. Parameters are estimated using the training set, and each observation in the test set is assigned to the component to which it has the highest membership probability. Parameter estimation for model-based classification and discriminant analysis proceeds in a similar fashion to model-based clustering; cf. McNicholas (2010).

4 Illustrations

4.1 Implementation and Evaluation

In the following illustrations, we fitted the MGHDs, the MMSGHDs, the McMSGHDs, and the MCGHDs using the corresponding functions available in the `MixGHD` package (Tortora et al., 2015) for `R` (R Core Team, 2014). We use k-means clustering to initialize the τ_{ig}. The adjusted Rand index (ARI; Hubert and Arabie, 1985) is used to compare predicted classifications with true classes. The ARI corrects the Rand index (Rand, 1971) for chance, its expected value under random classification is 0, and it takes a value of 1 when there is perfect class agreement. Steinley (2004) gives guidelines for interpreting ARI values.

4.2 Comparing the MGHDs, the MMSGHDs, the McMSGHDs, and the MCGHDs

To assess the classification performance of the MGHDs, the MMSGHDs, the McMSGHDs, and the MCGHDs, we consider four real data sets that are commonly used within the model-based clustering literature (Table 1).

We fitted the four mixtures to the real data sets in Table 1, with the number of components, G, set equal to the true number of classes. Table 2 displays the classification performance for the best fitting mixtures. The MCGHD generally performs at least as well as the best of the other three approaches. Specifically, the MCGHD gives the best classification performance — either outright or jointly — for the banknote and bankruptcy data. The

Table 1: Summary details for four data sets commonly used within the model-based clustering literature.

	Classes	n	p	R package
Banknote	2	200	7	`mclust` (Fraley et al., 2014)
Bankruptcy	2	66	2	`MixGHD` (Tortora et al., 2015)
HSCT	4	9780	4	Not available online
AIS	2	202	11	`EMMIXuskew` (Lee and McLachlan, 2013a)

	Original source
Banknote	Flury and Riedwyl (1988)
Bankruptcy	Alman (1968)
HSCT	Terry Fox Lab
AIS	Cook and Weisberg (1994)

MCGHD approach gives similar classification performance to the other approaches for the AIS data. For the HSCT data, the MCGHD outperforms two of the other three approaches and gives somewhat inferior performance compared to MGHD. Interestingly, the MGHD gives very good classification performance on three of the four data sets; however, this must be taken in context with its very poor classification performance on the bankruptcy data set (ARI ≈ 0). It is also interesting to compare the classification performance of the McMSGHD to the MMSGHD as well as that of the MGHD to the MMSGHD. The McMSGHD approach outperforms MMSGHD for two of the four data sets, giving the same performance on the other two. Interestingly, the MGHD approach outperforms the MMSGHD model on two of the four data sets. This highlights the fact that a mixture of multiple scaled distributions may well not outperform its single scaled analogue, and underlines the need for an approach with both the MSGHD and MGHD models as special cases. Finally, the results for the bankruptcy data illustrate that the MCGHD approach can give very good classification performance in situations where neither the MGHD nor the MMSGHD perform well.

Table 2: ARI values for the MCGHD, MGHD, MMSGHD, and McMSGHD approaches on four real data sets.

	G	MCGHD	MGHD	MMSGHD	McMSGHD
Banknote	2	0.980	0.980	0.980	0.980
Bankruptcy	2	0.824	0.052	0.198	0.256
HSCT	4	0.781	0.936	0.210	0.732
AIS	2	0.847	0.884	0.811	0.811

In this analysis, we took the number of components to be known. However, this is not realistic for clustering in general. Therefore, the analysis was repeated without this assumption. For each data set, all four approaches were run for values of G ranging from

one to a number two greater than the true number of classes, e.g., for the banknote data, we used $G = 1, \ldots, 4$. The results are identical to those given in Table 2, with the following exception: for the HSCT data, the MGHD approach led to a $G = 5$ component solution (ARI=0.660). Accordingly, when G is not fixed to the true value, the MCGHD gives the best classification performance — either outright or jointly — for three of the four data sets and performs similarly to the other approaches for the fourth (AIS).

Before moving to the next section, it is interesting to consider a contour plot of the MCGHD result for the bankruptcy data (Figure 4), which presents another example of the unusual component shapes that can be accommodated by our MCGHD model. The poor performance of the MGHD approach on these data can be explained by effect of the two outlying points on a single-scaled distribution.

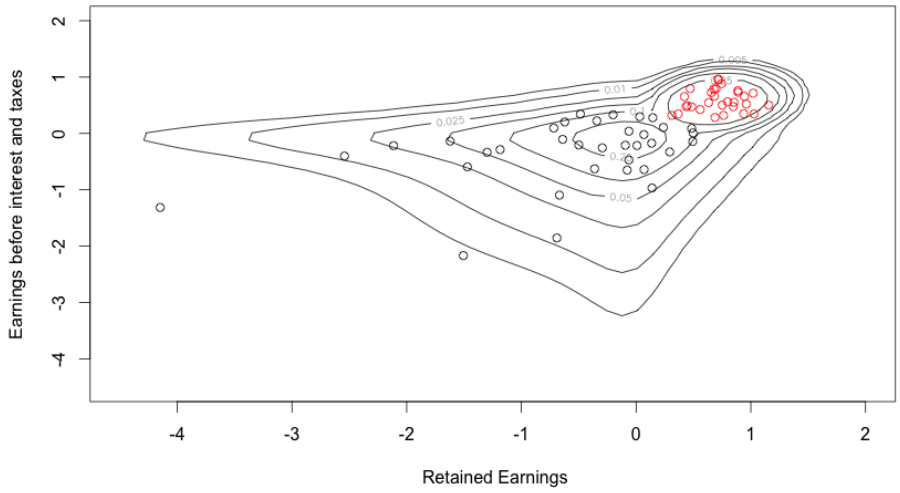

Figure 4: Contour plot of the best fitting MCGHD for the bankruptcy data, with colour denoting predicted classifications.

4.3 Comparison with MuST

Recently, several mixtures of skewed distributions have been proposed for model-based clustering and classification. Among them, mixtures of unrestricted skew-t distributions (MuST) and mixtures of restricted skew-t distributions (MrST) have been widely discussed in the literature. The terminology used to define these distributions, i.e., restricted and unrestricted, is questionable (cf. Azzalini et al., 2014); however, we will use this terminology herein to be consistent with Lee and McLachlan (2014). Lee and McLachlan (2013b, 2014) compare these methods using real data sets and find that, in general, the MuST outperforms MrST, skew-normal mixtures, and some other mixtures of skewed distributions. Accordingly, we choose

to compare our MCGHDs to the MuST approach, available in the `EMMIXuskew` package (Lee and McLachlan, 2013a) for R.

4.3.1 Comparing efficiency

To compare the computational times of the approaches[1], we used the `Colon` data set from the R package `plsgenomics` (Boulesteix et al., 2011), which contains 62 tissues (40 tumour tissues and 22 normal tissues) with 2000 genes each. We fitted one- and two-component MCGHDs, using only the tumour tissues when $G = 1$. To compare times, we selected subsets of the `Colon` data with $p = 2, 5, 10, 25, 50, 100$ dimensions, respectively. Figure 5 shows the average elapsed time (in seconds) for 25 replications of 100 iterations of each (G)EM algorithm for each value of p. From Figure 5, we can see that the average elapsed time for the MCGHDs increases linearly with the number of variables; however, the average elapsed time for the MuST increases exponentially with the number of variables. Relative to the MSGHD approach, the MuST approach is painfully slow, e.g., when $p = 5$, the elapsed time using MuST is 610.038 seconds for $G = 1$ and 1046.463 seconds for $G = 2$, and when $p = 10$ and $G = 2$, 100 EM iterations required more than 9 hours. Because of the excessive amounts of time that would be involved, we do not fit the MuST approach for $p > 10$ — this is the reason that the plots in Figure 5 have a different range of values on their respective x-axes.

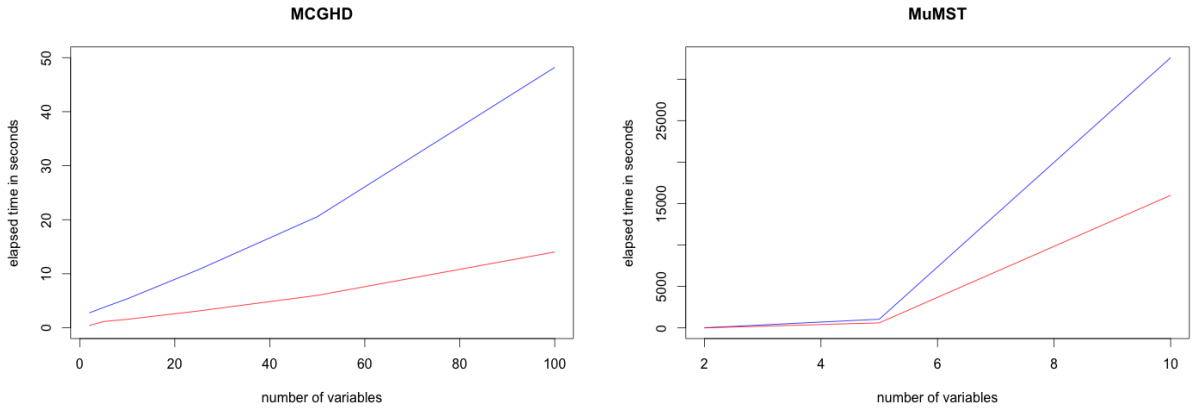

Figure 5: The average elapsed time to preform 100 (G)EM iterations when varying the number of variables for one- (red line) and two- (blue line) component MCGHD and MuST models.

[1]All code was run in R version 3.0.2 on Mac OS X version 10.6.8, with a 3.06 GHz Intel Core 2 Duo processor and 4 GB 800 MHz memory.

4.3.2 Comparing classification performance using real data

We use the real data sets in Table 1 to compare the classification performance of the MCGHD and MuST approaches. Both the MCGHDs and MuST were fitted to each data set for $G = 2, \ldots, 5$ components and the best fitting mixture was chosen using the BIC. For the HSCT data set we only fitted $G = 4$ component mixtures, following Lee and McLachlan (2013b). Furthermore, results herein may differ from those given by Lee and McLachlan (2013a) because we scaled each data set prior to analysis and used k-means clustering results to initialize the algorithms. The performance of the two approaches is summarized in Table 3. For each of the four data sets, the BIC selects MCGHDs with the correct number of components. The correct number of components is selected for the MuST approach for three of the four data sets. The MCGHD gives better classification performance than the MuMST on two of the four data sets, and *vice versa*.

Table 3: Number of components (G) and ARI values for the selected MCGHDs and MuST approaches for four real data sets.

	MCGHD		MuST	
	G	ARI	G	ARI
Banknote	2	0.980	3	0.851
Bankruptcy	2	0.824	2	0.549
HSCT	4	0.781	4	0.977
AIS	2	0.847	2	0.914

4.4 Model-Based Classification and Discriminant Analysis

Two real data sets that are popular in model-based classification and/or discriminant analysis applications (Table 1) are used to compare the classification performances of the MGHD, MSGHD, MCGHD, and cMSGHD approaches. We also compare these methods with Gaussian mixtures, as implemented for model-based discriminant analysis via the `mclust` package (Fraley et al., 2014) for R and for model-based classification via the `mixture` package (Browne and McNicholas, 2015).

Table 4: Summary details for two data sets commonly used within the model-based classification and/or discriminant analysis literature.

	Classes	n	p	R package	Original source
Bankruptcy	2	66	2	`MixGHD` (Tortora et al., 2015)	Alman (1968)
Diabetes	3	145	3	`mclust` (Fraley et al., 2014)	Reaven and Miller (1979)

To illustrate model-based classification and discriminant analysis, 25% of the observations in the bankruptcy data set and 27% of the observations in diabetes data are randomly designated unlabelled. For discriminant analysis purposes, these unlabelled observations form the test sets. Looking at the ARI values for the predicted classifications for unlabelled observations versus true classifications (Table 5), it is clear that the MCGHD gives comparable or superior performance to the other approaches. The model-based classification analysis of the bankruptcy data is particularly interesting because the approaches introduced herein have succeeded where Gaussian mixtures failed and because the MCGHD outperforms both of its special cases (MMSGHD and MGHD). The latter observation points to the fact that the MCGHD can perform well in situations where neither MMSGHD nor MGHD do well. The same phenomenon is observed for the model-based discriminant analysis of the bankruptcy data. For the diabetes data, all five approaches give similar performance under model-based classification and model-based discriminant analysis.

Table 5: ARI values, for unlabelled observations, associated with the Gaussian mixture (GM), MCGHD, MGHD, MMSGHD, and McMSGHD analyses of two real data sets.

		GM	MCGHD	MGHD	MMSGHD	McMSGHD
Classification	Bankruptcy	0.200	0.750	0.533	0.533	0.533
	Diabetes	0.805	0.856	0.805	0.941	0.941
Discriminant analysis	Bankruptcy	0.350	0.750	0.533	0.533	1.000
	Diabetes	0.805	0.856	0.856	0.856	0.856

5 Simulation study

We designed a simulation study to assess the classification abilities of our MCGHDs, MGHDs, MMSGHDs, McMSGHD, and the MuST on data generated from a multivariate normal distribution, a multivariate skew-normal distribution (SN), a GHD, and a MSGHD. In total, we consider 36 scenarios. Two-, three-, and four-component mixtures were generated with 200 p-dimensional vectors each. Furthermore, every component is centred on a different point with the centres uniformly distributed in a hypercube (side length 50). The $p \times p$ diagonal matrix $\boldsymbol{\Sigma}_g$ is randomly generated with off diagonal elements in the interval $[0, 0.6]$ and $\text{diag}(\boldsymbol{\Sigma}_g) = 1$. The skewness parameter $\boldsymbol{\alpha}_g$ is randomly generated in the interval $[-6, 6]$, and the values of the other parameters are $\boldsymbol{\omega}_g = \mathbf{1}_p,\ \boldsymbol{\lambda}_g = -0.5\mathbf{1}_p,\ \omega_0 = 1,$ and $\lambda_0 = -0.5$. The normally distributed data sets were generated using the R function `rnorm`, the skew normal data sets were generated using the `rdmsn` function from the R package `EMMIXskew`, and the generalized hyperbolic and multiple scaled generalized hyperbolic data were generated using the stochastic relationships given in (14) and (17), respectively. Because fitting MuST is very time consuming, we only apply it for $p = 5$. For each scenario we generate 50 data sets, and

we apply every method to each of them. We measure the ARI for every result, in Table 6 we report the 5th, 50th, and 95th percentile of the ARI distribution for each scenario. Results for the MuST are shown in Table 7, where we also show the average elapsed time for 50 replications of the algorithm[2].

5.1 Simulation Results

Table 6 shows the 5th, 50th, and 95th percentile of the ARI distribution for each method with $G = \{2, 3, 4\}$ and $p = \{5, 10, 50\}$. For the data sets generated from a multivariate Gaussian distribution the MCGHDs, MMSGHDs, and McMSGHDs give excellent results, median equal to 1 in every scenario. The MGHDs has a lower median when $G = 4$, while the MuST has a lower median when $G = 3$ and $G = 4$. For the data sets generated from the skew normal distribution, the MCGHDs outperform the other methods with a median equal to 1 in every scenario. The MGHDs and MMSGHDs have a lower median when $p = 50$ and $G = 4$, while McMSGHD has a lower median when $G = 2$, and the MuST has a lower median when $G = 4$. When the data are generated using a GHD all the methods give good results with the exception of the scenario with $p = 5$ and $G = 4$, characterized by lower median values for the MCGHDs, MMSGHDs, and McMSGHDs. On the data sets generated using a MSGHD all the methods perform well, with the exception of the MuST which has a median of 0.69 when $G = 4$.

6 Discussion and Conclusion

Novel MCGHDs, MMSGHDs, and McMSGHDs models have been introduced and applied for model-based clustering, classification, and discriminant analysis. The GHD is a flexible distribution, capable of handling skewness and heavy tails, and has many well known distributions as special or limiting cases. Furthermore, it is a normal variance-mean mixture, arising via a relationship between a multivariate Gaussian and an univariate GIG distribution. The MSGHD extends the GHD to include a multivariate GIG distribution, increasing the flexibility of the model. However, the GHD is not a special case of the MSGHD; hence, we created MCGHDs, which has both the GHD and MSGHD as special cases.

The McMSGHD approach was introduced as a convex version of the MMSGHD. This point deserves some further discussion. The extension of the multivariate-t distribution to multiple scale was carried out by Forbes and Wraith (2014), the multiple scaled multivariate t-distribution cannot be quasi-concave, i.e., the clusters associated with a mixture of multiple scaled multivariate t-distributions cannot be convex. We have seen examples where the MMSGHD can put multiple clusters into one component, and the McMSGHD has an important role in helping to prevent this; if both approaches are fitted and lead to different numbers of components, then further attention is deserved.

[2]Using a 32-core Intel Xeon E5 server with 256GB RAM running 64-bit CentOS.

Table 6: ARI percentiles for the MCGHDs, MGHDs, MMSGHDs, and McMSGHDs fitted to the simulated data sets.

Distribution	p	G	MCGHD			MGHD			MMSGHD			McMSGHD		
			.05	.50	.95	.05	.50	.95	.05	.50	.95	.05	.50	.95
$\mathcal{N}$	5	2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\mathcal{N}$	10	2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\mathcal{N}$	50	2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\mathcal{N}$	5	3	1.00	1.00	1.00	0.44	1.00	1.00	0.45	1.00	1.00	0.56	1.00	1.00
$\mathcal{N}$	10	3	1.00	1.00	1.00	0.44	1.00	1.00	0.46	1.00	1.00	1.00	1.00	1.00
$\mathcal{N}$	50	3	0.44	1.00	1.00	0.44	1.00	1.00	0.44	1.00	1.00	0.45	1.00	1.00
$\mathcal{N}$	5	4	0.64	1.00	1.00	0.63	0.66	1.00	0.64	1.00	1.00	0.63	1.00	1.00
$\mathcal{N}$	10	4	0.63	1.00	1.00	0.63	0.64	1.00	0.00	1.00	1.00	0.63	1.00	1.00
$\mathcal{N}$	50	4	1.00	1.00	1.00	0.63	0.63	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\mathcal{SN}$	5	2	0.00	1.00	1.00	0.00	1.00	1.00	0.00	1.00	1.00	0.00	0.21	1.00
$\mathcal{SN}$	10	2	0.00	1.00	1.00	1.00	1.00	1.00	0.07	1.00	1.00	0.00	0.31	1.00
$\mathcal{SN}$	50	2	0.00	1.00	1.00	0.45	1.00	1.00	0.00	1.00	1.00	0.00	0.68	1.00
$\mathcal{SN}$	5	3	0.57	1.00	1.00	0.44	1.00	1.00	0.54	1.00	1.00	0.45	1.00	1.00
$\mathcal{SN}$	10	3	0.57	1.00	1.00	0.44	1.00	1.00	0.55	1.00	1.00	0.55	1.00	1.00
$\mathcal{SN}$	50	3	0.51	1.00	1.00	0.00	1.00	1.00	0.00	1.00	1.00	0.28	1.00	1.00
$\mathcal{SN}$	5	4	0.57	1.00	1.00	1.00	1.00	1.00	0.54	1.00	1.00	0.45	1.00	1.00
$\mathcal{SN}$	10	4	0.57	1.00	1.00	1.00	1.00	1.00	0.54	1.00	1.00	0.45	1.00	1.00
$\mathcal{SN}$	50	4	0.00	1.00	1.00	0.00	0.64	1.00	0.00	0.71	1.00	0.00	0.70	1.00
$\mathcal{GHD}$	5	2	0.03	1.00	1.00	1.00	1.00	1.00	0.14	1.00	1.00	0.04	1.00	1.00
$\mathcal{GHD}$	10	2	0.13	1.00	1.00	1.00	1.00	1.00	0.98	1.00	1.00	0.04	1.00	1.00
$\mathcal{GHD}$	50	2	0.86	1.00	1.00	0.71	1.00	1.00	0.44	1.00	1.00	0.44	1.00	1.00
$\mathcal{GHD}$	5	3	0.42	0.90	1.00	0.50	1.00	1.00	0.44	0.97	1.00	0.44	1.00	1.00
$\mathcal{GHD}$	10	3	0.27	1.00	1.00	0.50	1.00	1.00	0.48	1.00	1.00	0.44	1.00	1.00
$\mathcal{GHD}$	50	3	0.42	1.00	1.00	0.50	1.00	1.00	0.48	1.00	1.00	0.44	1.00	1.00
$\mathcal{GHD}$	5	4	0.36	0.73	1.00	0.64	1.00	1.00	0.59	0.77	1.00	0.54	0.68	1.00
$\mathcal{GHD}$	10	4	0.48	0.97	1.00	0.66	1.00	1.00	0.58	1.00	1.00	0.62	1.00	1.00
$\mathcal{GHD}$	50	4	0.42	1.00	1.00	0.64	1.00	1.00	0.62	1.00	1.00	0.62	1.00	1.00
$\mathcal{MSGHD}$	5	2	0.88	1.00	1.00	0.99	1.00	1.00	0.79	1.00	1.00	0.46	1.00	1.00
$\mathcal{MSGHD}$	10	2	0.97	1.00	1.00	1.00	1.00	1.00	0.97	1.00	1.00	0.66	1.00	1.00
$\mathcal{MSGHD}$	50	2	0.94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\mathcal{MSGHD}$	5	3	0.81	0.98	1.00	0.45	1.00	1.00	0.51	0.98	1.00	0.52	0.99	1.00
$\mathcal{MSGHD}$	10	3	0.83	1.00	1.00	0.44	1.00	1.00	0.45	1.00	1.00	0.45	1.00	1.00
$\mathcal{MSGHD}$	50	3	0.44	1.00	1.00	0.44	1.00	1.00	0.45	1.00	1.00	1.00	1.00	1.00
$\mathcal{MSGHD}$	5	4	0.64	0.98	1.00	0.62	0.99	1.00	0.65	0.96	1.00	0.62	0.98	1.00
$\mathcal{MSGHD}$	10	4	0.62	1.00	1.00	0.63	1.00	1.00	0.63	1.00	1.00	0.63	1.00	1.00
$\mathcal{MSGHD}$	50	4	0.61	0.99	1.00	0.63	1.00	1.00	0.62	1.00	1.00	0.63	1.00	1.00

Comparing the MGHD, MMSGHD, McMSGHD, and MCGHD approaches yielded some interesting results. Amongst them, we see that the MMSGHD does not necessarily outperform the MGHD; far from it, in fact, because the MGHD approach gave better clustering

Table 7: ARI percentiles for the MuST fitted to the simulated data sets.

Distribution	p	G	MuST			Average elapsed time
			.05	.50	.95	
$\mathcal{N}$	5	2	1.00	1.00	1.00	$202s$
$\mathcal{N}$	5	3	0.45	0.52	1.00	$5294.213s \approx 1.5h$
$\mathcal{N}$	5	4	0.63	0.71	0.87	$7050.646s \approx 2h$
$\mathcal{SN}$	5	2	1.00	1.00	1.00	$8396.899s \approx 2h$
$\mathcal{SN}$	5	3	0.51	1.00	1.00	$12985.54s \approx 3.5h$
$\mathcal{SN}$	5	4	0.52	0.72	1.00	$23500.72s \approx 6.5h$
$\mathcal{GHD}$	5	2	1.00	1.00	1.00	$5892.214s \approx 1.5h$
$\mathcal{GHD}$	5	3	1.00	1.00	1.00	$9686.482s \approx 2.5h$
$\mathcal{GHD}$	5	4	1.00	1.00	1.00	$13791.28s \approx 4h$
$\mathcal{MSGHD}$	5	2	1.00	1.00	1.00	$5313.021s \approx 1.5h$
$\mathcal{MSGHD}$	5	3	1.00	1.00	1.00	$9140.247s \approx 2.5h$
$\mathcal{MSGHD}$	5	4	0.63	0.69	1.00	$13247.85s \approx 3.5h$

performance than the MMSGHD approach on two of the four real data sets we considered, as well as identical performance on a third. This underlines the fact that a mixture of multiple scaled distributions may well not outperform its single scaled analogue, and highlights the benefit of approaches with both a multiple scaled distribution and its single scaled analogue as special cases. The MCGHDs represent one such approach, with the MSGHD and MGHD models as special cases. The approaches introduced herein, as well as the MGHDs, have been made publicly available via the `MixGHD` package for R.

References

Aitken, A. (1926). On Bernoulli's numerical solution of algebraic equations. *Proceedings of the Royal Society of Edimburgh 46*, 289–305.

Alman, E. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankrupcy. *J Finance 23*(4), 589–609.

Azzalini, A., R. P. Browne, M. G. Genton, and P. D. McNicholas (2014). Comparing two formulations of skew distributions with special reference to model-based clustering. *arXiv preprint*.

Baek, J., G. J. McLachlan, and L. Flack (2010). Mixtures of factor analyzers with common factor loadings: Applications to the clustering and visualization of high-dimensional data. *IEEE Transactions on Pattern Analysis and Machine Intelligence 32*(7), 1298–1309.

Barndorff-Nielsen, O. (1978). Hyperbolic distributions and distributions on hyperbolae. *Scandinavian Journal of statistics*, 151–157.

Barndorff-Nielsen, O., J. Kent, and M. Sørensen (1982). Normal variance-mean mixtures and z distributions. *International Statistical Review / Revue Internationale de Statistique 50*(2), 145–159.

Böhning, D., E. Diez, R. Scheub, P. Schlattmann, and B. Lindsay (1994). The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family. *Annals of the Institute of Statistical Mathematics 46*, 373–388.

Boulesteix, A., S. Lambert-Lacroix, J. Peyre, and K. Strimmer (2011). *plsgenomic: PLS analyses for genomics*. R package version 1.2-6.

Bouveyron, C. and C. Brunet-Saumard (2014). Model-based clustering of high-dimensional data: A review. *Computational Statistics and Data Analysis 71*, 52–78.

Bouveyron, C., S. Girard, and C. Schmid (2007). High-dimensional data clustering. *Computational Statistics & Data Analysis 52*(1), 502–519.

Browne, R. P. and P. D. McNicholas (2014). Estimating common principal components in high dimensions. *Advances in Data Analysis and Classification 8*(2), 217–226.

Browne, R. P. and P. D. McNicholas (2015). A mixture of generalized hyperbolic distributions. *Canadian Journal of Statistics to appear*.

Browne, R. P., P. D. McNicholas, and M. D. Sparling (2012). Model-based learning using a mixture of mixtures of Gaussian and uniform distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence 34*(4), 814–817.

Campbell, J. G., F. Fraley, F. Murtagh, and A. E. Raftery (1997). Linear flaw detection in woven textiles using model-based clustering. *Pattern Recognition Letters 18*(14), 1539–1548.

Celeux, G. and G. Govaert (1995). Gaussian parsimonious clustering models. *Pattern Recognition 28*(5), 781–793.

Cook, R. D. and S. Weisberg (1994). *An Introduction to Regression Graphics*. John Wiley & Sons, New York.

Dasgupta, A. and A. E. Raftery (1998). Detecting features in spatial point processed with clutter via model-based clustering. *Journal of American Statistical Association 93*(441), 294–302.

Debreu, G. and T. C. Koopmans (1982). Additively decomposed quasiconvex functions. *Mathematical Programming 24*(1), 1–38.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B 39*(1), 1–38.

Eltoft, T., T. Kim, and T.-W. Lee (2006). Multivariate scale mixture of gaussians modeling. In *Independent Component Analysis and Blind Signal Separation*, pp. 799–806. Springer.

Flury, B. and H. Riedwyl (1988). *Multivariate Statistics: A practical approach*. London: Chapman & Hall.

Forbes, F. and D. Wraith (2014). A new family of multivariate heavy-tailed distributions with variable marginal amounts of tailweights: Application to robust clustering. *Statistics and Computing 24*(6), 971–984.

Fraley, C., A. Raftery, and L. Scrucca (2014). *mclust: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation*. version 4.3.

Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association 97*(458), 611–631.

Franczak, B. C., R. P. Browne, and P. D. McNicholas (2014). Mixtures of shifted asymmetric Laplace distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence 36*(6), 1149–1157.

Franczak, B. C., C. Tortora, R. P. Browne, and P. D. McNicholas (2015). Unsupervised learning via mixtures of skewed distributions with hypercube contours. *Pattern Recognition Letters 58*(1), 69 – 76.

Gneiting, T. (1997). Normal scale mixtures and dual probability densities. *Journal of Statistical Computation and Simulation 59*(4), 375–384.

Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of Classification 2*(1), 193–218.

Hunter, D. R. and K. Lange (2000). Quantile regression via an MM algorithm. *Journal of Computational and Graphical Statistics 9*(1), 60–77.

Johnson, N., S. Kotz, and N. Balakrishnan (1994). *Continuous Univariate Distributions*. New York: John Wiley & Sons.

Jørgensen, B. (1982). *Statistical Properties of the Generalized Inverse Gaussian Distribution*. New York: Springer-Verlag.

Karlis, D. and L. Meligkotsidou (2007). Finite mixtures of multivariate Poisson distributions with application. *Journal of Statistical Planning and Inference 137*(6), 1942–1960.

Karlis, D. and A. Santourian (2009). Model-based clustering with non-elliptically contoured distributions. *Statistics and Computing 19*(1), 73–83.

Kiers, H. A. (2002). Setting up alternating least squares and iterative majorization algorithms for solving various matrix optimization problems. *Computational Statistics and Data Analysis 41*(1), 157–170.

Kotz, S., T. J. Kozubowski, and K. Podgorski (2001). *The Laplace Distribution and Generalizations: A Revisit with Applications to Communications, Economics, Engineering, and Finance* (1st ed.). Burkhauser Boston.

Kotz, S. and S. Nadarajah (2004). *Multivariate t-distributions and their applications*. Cambridge University Press.

Lee, S. X. and G. J. McLachlan (2013a). *EMMIXuskew: Fitting Unrestricted Multivariate Skew t Mixture Models*. R package version 0.11-5.

Lee, S. X. and G. J. McLachlan (2013b). On mixtures of skew normal and skew t-distributions. *Advances in Data Analysis and Classification 7*(3), 241–266.

Lee, S. X. and G. J. McLachlan (2014). Finite mixtures of multivariate skew t-distributions: some recent and new results. *Statistics and Computing 24*(2), 181–202.

Lin, T. I. (2009). Maximum likelihood estimation for multivariate skew normal mixture models. *Journal of Multivariate Analysis 100*(2), 257–265.

Lin, T. I. (2010). Robust mixture modeling using multivariate skew t distributions. *Statistics and Computing 20*(3), 343–356.

Lindsay, B. (1995). Mixture models: Theory, geometry and applications. In *NSF-CBMS Regional Conference Series in Probability and Statistics*, Volume 5, California: Institute of Mathematical Statistics: Hayward.

McLachlan, G. J. and D. Peel (2000). *Finite Mixture Models*. Wiley Interscience, New York.

McLachlan, G. J., D. Peel, and R. W. Bean (2003). Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics and Data Analysis 41*(3), 379–388.

McNeil, A. J., R. Frey, and P. Embrechts (2005). *Quantitative risk management: concepts, techniques and tools*. Princeton university press.

McNicholas, P. D. (2010). Model-based classification using latent Gaussian mixture models. *Journal of Statistical Planning and Inference 140*(5), 1175–1181.

McNicholas, P. D. and T. Murphy (2008). Parsimonious Gaussian mixture models. *Statistics and Computing 18*(3), 285–296.

McNicholas, P. D. and T. Murphy (2010). Model-based clustering of microarray expression data via latent Gaussian mixture models. *Bioinformatics 26*(21), 2705–2712.

Montanari, A. and C. Viroli (2011). Maximum likelihood estimation of mixtures of factor analyzers. *Computational Statistics and Data Analysis 55*(9), 2712–2723.

Murray, P. M., R. B. Browne, and P. D. McNicholas (2014). Mixtures of skew-t factor analyzers. *Computational Statistics and Data Analysis 77*, 326–335.

Niculescu, C. and L. Persson (2006). *Convex Functions and Their Applications*. New York: Springer.

Ortega, J. M. and W. C. Rheinboldt (1970). *Iterative Solutions of Nonlinear Equations in Several Variables*. New York: Academic Press.

Peel, D. and G. J. McLachlan (2000). Robust mixture modelling using the t distribution. *Statistics and Computing 10*(4), 339–348.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association 66*(336), 846–850.

Reaven, G. and R. Miller (1979). An attempt to define the nature of chemical diabetes using a multidimensional analysis. *Diabetologica 16*, 17–24.

Rockafellar, R. T. and R. J. B. Wets (2009). *Variational Analysis*. New York: Springer-Verlag.

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics 6*(2), 461–464.

Steinley, D. (2004). Properties of the Hubert-Arable adjusted Rand index. *Psychological methods 9*(3), 386.

Titterington, D. M., A. F. M. Smith, and U. E. Makov (1985). *Statistical Analysis of Finite Mixture Distributions*. Chichester: John Wiley & Sons.

Tortora, C., R. P. Browne, B. C. Franczak, and P. D. McNicholas (2015). MixGHD: Model based clustering and classification using the mixture of generalized hyperbolic distributions. *R package version 1.5*.

Tortora, C., P. D. McNicholas, and R. P. Browne (2015). A mixture of generalized hyperbolic factor analyzers. *Advanced in Data Analysis and Classification In press*.

Vrbik, I. and P. D. McNicholas (2012). Analytic calculations for the EM algorithm for multivariate skew-mixture models. *Statistics and Probability Letters 82*(6), 1169–1174.

Vrbik, I. and P. D. McNicholas (2014). Parsimonious skew mixture models for model-based clustering and classification. *Computational Statistics and Data Analysis 71*, 196–210.

Wraith, D. and F. Forbes (2015). Clustering using skewed multivariate heavy tailed distributions with flexible tail behaviour. arXiv:1408.0711.